arXiv:1509.08830v1 [cs.LG] 29 Sep 2015

How to Formulate and Solve Statistical Recognition and Learning Problems

M.I.Schlesinger, schles@irtc.org.ua
International Research and Training Center of
Informational Technologies and Systems
Ukrainian Academy of Sciences
Kiev

E.V.Vodolazskiy, waterlaz@gmail.com
International Research and Training Center of
Informational Technologies and Systems
Ukrainian Academy of Sciences
Kiev

September 30, 2015

Abstract

We formulate problems of statistical recognition and learning in a common framework of complex hypothesis testing. Based on arguments from multi-criteria optimization, we identify strategies that are improper for solving these problems and derive a common form of the remaining strategies. We show that some widely used approaches to recognition and learning are improper in this sense. We then propose a generalized formulation of the recognition and learning problem which embraces the whole range of sizes of the learning sample, including the zero size. Learning becomes a special case of recognition without learning. We define the concept of closest to optimal strategy, being a solution to the formulated problem, and describe a technique for finding such a strategy. On several illustrative cases, the strategy is shown to be superior to the widely used learning methods based on maximal likelihood estimation.

Keywords: complex object recognition, learning, multi-criteria optimization, Bayesian strategy, small sample problem

1 Introduction

We consider three wide directions in theoretical and applied statistics: complex hypothesis testing [5, 11, 13, 14], the empirical Bayesian approach [4, 15, 16, 17] and learning in pattern recognition [7, 22]. Each of these fields is well-known and the list of relevant publications would be longer than this paper.

However, the methods developed in the three fields are not fully consistent with one another. Indeed, the most popular learning methods seem to have been developed independently of complex hypothesis testing and sometimes even contradict it. It seems as if the well-known ideas of self-learning or unsupervised learning [6, 19, 20] appeared independently of the empirical Bayesian approach. Such an isolated existence of learning methods results in their rather fragile foundation.

Learning in pattern recognition deals with the situation when the statistical model of an object to be recognized is not defined uniquely and only a set of models is known that includes the true model. The common feature of modern learning methods is that a so-called learning sample is given and based on this sample a certain model is selected from the given set. Then the recognition strategy is derived as if the selected model was the true one. Usually, the selected model is a consistent estimate of the true model and therefore this approach is acceptable, provided that the learning sample is large enough. However, if the learning sample has limited size the approach gives no guarantee for subsequent recognition. The situation is known as the "small sample problem".

Thus there is an obvious gap in our knowledge of statistical recognition under uncertain statistical model. Strictly speaking, satisfactory clarity is achieved only in two extreme cases. When no learning sample is available the methods for complex hypothesis testing are applicable, such as the minimax approach. When a large enough learning sample is available the learning methods are applicable. Practical situations fall in the gap between these two extreme cases because any sample has a certain size, it is not arbitrarily large. The situation is discussed in more detail in [19, page 272].

The paper fills this gap. Using the concepts of multi-criteria optimization [8, 9], we analyze the empirical Bayesian approach, learning in pattern recognition and complex hypothesis testing in a unified framework. Based on this unification, we formulate the learning problem as a special case of recognition without learning, the latter in turn being a special case of complex hypothesis testing. The proposed formulation covers the whole range of learning sample sizes, including zero size. Thus, the small sample problem disappears as an independent problem and is embraced by the proposed formulation.

The paper is organized as follows:

Section 2. Recognition without learning is considered under an uncertain statistical model of an object to be recognized. The concepts of improper and Bayesian strategies are defined and the dichotomy theorem is proved that each strategy is either Bayesian or improper. It is shown that the strategies based on maximum likelihood model estimate are improper. The concept of closest to optimal recognition is defined.

Section 3. The seminal ideas of H.Robbins [16], which addressed recognition of compound objects and initiated the empirical Bayesian approach [15], are revisited. Robbins' ideas are formalized as finding a closest to optimal strategy. It is shown that strategies based on the well-known EM-algorithms differ from the closest to optimal ones and are improper for some objects.

Section 4. Recognition with learning is formulated as a special case of

recognition without learning. It is shown that the widely used strategies based on maximal likelihood model estimate are improper. It is also shown that the minimax approach to statistical decision making is unsuitable for solving the learning problem. The concept of closest to optimal learning is defined and its formal properties are analyzed.

Section 5. The technique of closest to optimal recognition is described.

Section 6. Closest to optimal learning procedures are tested on several illustrative examples. The tests include supervised and unsupervised learning. The test with supervised learning shows that the closest to optimal learning procedure is superior to the procedure based on maximum-likelihood model estimation. The tests with unsupervised learning compare the closest to optimal procedures with EM-algorithms and some heuristic methods. The proposed approach yields the best result.

Section 7. The results of the paper as well as their consequences are summarized.

The results of Sections 2, 3 and 4 are mostly negative, arguing that the commonly used approaches have fundamental drawbacks and thus new approaches are needed. The positive results concern mainly the concept of closest to optimal learning and are described in Sections 4, 5 and 6.

2 Complex object recognition

The main concepts of complex hypothesis discrimination have been formulated in the general statistical decision theory [5, 11, 13, 14]. Our aim is to relate the result of this theory to statistical recognition and machine learning [7, 22]. Therefore we formulate the main concepts of statistical decision theory using pattern recognition terminology. Basic concepts of our consideration are a set of observable signals, a set of hidden states and a set of models. All three sets are assumed to be finite. This simplification allows to articulate the main ideas of the article using the simplest mathematical tools. Allowing some of the sets to be infinite might obscure the main ideas by more fine mathematical formulations. Nevertheless, some examples below involve intervals of real numbers. Our theoretical results are applied to such examples without further ado. Probability density is simply used instead of probability.

Let an object to be recognized be characterized by two parameters x and y that take values from finite sets X and Y respectively. The parameter x is an observable signal generated by the object. The parameter y is a hidden state of the object. The signal x and the state y form a random pair, so that a probability distribution $p : X \times Y \to \mathbb{R}$ over the set of pairs (x, y) exists. However, the probability distribution p may be either known or not known and this makes the distinction between simple and complex objects.

Definition 1. *A simple object is represented by a triple* $\langle X, Y, p : X \times Y \to \mathbb{R} \rangle$ *where* X *is a set of signals,* Y *is a set of states,* $p(x, y)$ *is a joint probability of signal* $x \in X$ *and state* $y \in Y$*.*

We consider more complex cases when the probability distribution p is not known and only a class of distributions is known that contains p. This uncertainty is represented by a finite set Θ of models θ, so that the probability of a pair (x, y) is defined for each model θ by $p(x, y; \theta)$.

Definition 2. *A complex object is represented by a quadruple* $\langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle$ *where* X *is a set of signals,* Y *is a set of states,* Θ *is a set of models and* $p(x, y; \theta)$ *is a joint probability of signal* $x \in X$ *and state* $y \in Y$ *for the model* $\theta \in \Theta$.

The above-defined complex object is not to be confused with what can be called a *pseudo-complex object*. The latter takes place when the model θ is a random variable with a known a priori probability distribution $p_\Theta \colon \Theta \to \mathbb{R}$. In fact, a pseudo-complex object does not differ from a simple one because for such an object the joint probability

$$p(x, y) = \sum_{\theta \in \Theta} p_\Theta(\theta)\, p(x, y; \theta)$$

is defined for every pair $x \in X$, $y \in Y$. In a complex object the model θ is not random but fixed although unknown. It is only known that the model belongs to a given set. This paper deals with complex objects, neither with simple ones nor with pseudo-complex ones.

Recognition means making a reasonable decision about the hidden state based on the observed signal. Below, we will also consider cases of the decision about current object state being based not only on the currently observed signal but on the base of additional information called learning information. In order to discriminate these two cases we use a term recognition without learning when the decision is based only on the observed signal and recognition with learning when the decision is based both on the signal and on the learning information.

Recognition without learning is formalized by a function $q_X \colon Y \times X \to \mathbb{R}$, called a randomized strategy. The values $q_X(y' \,|\, x), y' \in Y, x \in X$, of the function q_X are conditional probabilities. Let Q_X be the set of all possible strategies, that is, the set of functions $q_X \colon Y \times X \to \mathbb{R}$ satisfying $q_X(y' \,|\, x) \geq 0$ for every pair $y' \in Y$, $x \in X$ and $\sum\limits_{y' \in Y} q_X(y' \,|\, x) = 1$ for every $x \in X$. The strategy $q_X \colon Y \times X \to \mathbb{R}$ defines recognition as follows: if the signal x is observed then the decision "the object is in state y'" is made with probability $q_X(y' \,|\, x)$. The reasonableness of a strategy is formalized by a function $w \colon Y \times Y \to \mathbb{R}$, called a loss function. Its value $w(y, y')$, $y \in Y$, $y' \in Y$, is the loss of making a decision that the object state is y' when the true object state is y.

The statistical theory of complex hypothesis testing is an appropriate mathematical formalization for recognition of complex objects, not only simple ones. However, strategies that are widely used in recognition practice have almost nothing in common with this formalization.

Example 1. Let x be a picture and $y \in Y$ be the name of a letter drawn on the picture. The letters are drawn by some known person. A priori probabilities $p_Y(y), y \in Y$, of the letters are known for this person, as well as the conditional probabilities $p_{X \,|\, Y}(x \,|\, y)$. This is the case of a simple object.

Suppose now that there is a known finite team Θ of persons θ drawing pictures and the person θ that provides a picture for recognition is picked randomly with known probability $p_\Theta(\theta)$. It is necessary to recognize the letter drawn on the picture, not the person who drew it. One can assume that a priori probabilities $p_Y(y)$, $y \in Y$, of the letters do not depend on the person. However, the picture x depends both on the letter y and on the person $\theta \in \Theta$ who drew it. This dependency is given by conditional probabilities $p_{X\,|\,Y}(x\,|\,y;\theta)$, defined for every picture $x \in X$, every letter $y \in Y$ and every person $\theta \in \Theta$. This is the case of a pseudo-complex object.

In fact, this case does not differ from the case of a simple object. Indeed, one may substitute the whole team Θ of persons with a single abstract person θ^* who draws a random picture x of a letter y with conditional probability $p^*(x|y) = \sum\limits_{\theta\in\Theta} p_\Theta(\theta) p_{X\,|\,Y}(x\,|\,y;\theta)$.

Finally, let us consider the case when all pictures presented for recognition are drawn by a single person θ^0 from a known team Θ. However, this person is unknown. So, the model θ^0 is fixed, not random, and a concept of a priori probability distribution p_Θ becomes unappropriated in this case. Same as before, it is necessary to recognize the letter, not the person who drew it. This is the case of a complex object.

Commonly used algorithms for this task make the decisions

$$y^* = \operatorname*{argmax}_{y\in Y} \max_{\theta\in\Theta} p(x\,|\,y;\theta) \quad \text{or} \quad y^* = \operatorname*{argmax}_{y\in Y} \max_{\theta\in\Theta} p(x,y;\theta).$$

Neither of these two decisions follow from the statistical decision theory. Moreover, we will show later that both contradict this theory. ■

Statistical recognition problems for simple and complex objects are essentially different from each other in nature. In the case of a simple object, every strategy q_X is characterized by a single number, namely the expected value $R_X(q_X)$ of the losses,

$$R_X(q_X) = \sum_{x\in X} p(x) \sum_{y'\in Y} q_X(y'\,|\,x) \sum_{y\in Y} p(y\,|\,x)\, w(y,y'),$$

called the risk of strategy q_X. This leads to the known problem of Bayes risk (single-criterion) minimization over the strategy set. In the case of a complex object, every strategy q_X is characterized by $|\Theta|$ numbers $R_X(q_X,\theta)$, $\theta \in \Theta$, namely the risks of the strategy with respect to every model,

$$R_X(q_X,\theta) = \sum_{x\in X} \sum_{y'\in Y} \sum_{y\in Y} q_X(y'\,|\,x)\, p(y,x;\theta)\, w(y,y'). \tag{1}$$

The complex object case goes beyond the scope of single-criterion optimization and is subject to multi-criteria optimization. We follow its main ideas (see, for example, the book [8]) and translate them into the context of our problem.

First of all, we define some particularly bad strategies that do not optimize any reasonable criterion and thus they are useless.

Definition 3. *A strategy q'_X predominates a strategy q''_X if the strong inequality $R_X(q'_X, \theta) < R_X(q''_X, \theta)$ holds for every model $\theta \in \Theta$; a strategy q^0_X is called improper if a strategy q'_X exists that predominates q^0_X.*

One can see that the definition of improper strategy is rather strong. It does not embrace all strategies that can be treated as bad ones. For example, each strategy $\operatorname{argmin}_{q_X \in Q_X} R_X(q_X, \theta^*)$, $\theta^* \in \Theta$, is not improper but its usefulness is questionable because each such strategy takes into account only one single model θ^* and ignores all the others. Therefore, if some strategy is not improper it does not mean that it is unconditionally suitable for concrete application. However, if a strategy is proved to be improper it shows that it has a serious defect independently of the expected application.

We want to exclude all improper strategies from consideration and derive the common form of all other strategies. Such a dichotomy is possible due to the known result in multi-criteria optimization [8, Theorem 3.5]. The theorem should be central to complex object recognition both with and without learning. However, the known recognition methods have been developed and continue being developed as if the theorem did not exist. This is not surprising because the theorem has been formulated and proved in its own conceptual and terminological context, which is far from pattern recognition. We translate and even prove the theorem using the concepts of our article.

Let $\tau\colon \Theta \to \mathbb{R}$ be a function, called the weight function, satisfying $\tau(\theta) \geq 0$ for every $\theta \in \Theta$ and $\sum_{\theta \in \Theta} \tau(\theta) = 1$. The set of all such functions is denoted by T.

Definition 4. *The strategy $\operatorname{argmin}_{q_X \in Q_X} \sum_{\theta \in \Theta} \tau(\theta) R_X(q_X, \theta)$ is called Bayesian with respect to the weight function τ. A strategy q_X is called Bayesian if a weight function $\tau \in T$ exists with respect to which the strategy q_X is Bayesian.*

By this definition, every Bayesian strategy is optimal in a sense that it minimizes the Bayesian risk of recognizing a pseudo-complex object with a priori probability $\tau(\theta)$ of model θ. For a complex model, such treatment of Bayesian strategy is meaningless because the model is not random and its a priori probability distribution is not defined. However, the concept of Bayesian strategies allows to partition the set of all strategies into the set of improper strategies and the set of all other strategies. This dichotomy is given by the following theorem, which states that every strategy is either Bayesian or improper.

Theorem 1. *For every strategy $q^0_X \in Q_X$, either a weight function $\tau^* \in T$ exists such that*

$$q^0_X = \operatorname*{argmin}_{q_X \in Q_X} \sum_{\theta \in \Theta} \tau^*(\theta) R_X(q_X, \theta) \tag{2}$$

or a strategy q^*_X exists such that the inequality

$$R_X(q^*_X, \theta) < R_X(q^0_X, \theta) \tag{3}$$

holds for every $\theta \in \Theta$. These two properties of strategies are incompatible.

Proof. Let the function $F\colon T \times Q_X \to \mathbb{R}$ be defined by

$$F(\tau, q_X) = \sum_{\theta \in \Theta} \tau(\theta)[R_X(q_X, \theta) - R_X(q_X^0, \theta)].$$

Let the strategy q_X^* and the weight function τ^* be defined by

$$q_X^* = \operatorname*{argmin}_{q_X \in Q_X} \max_{\tau \in T} F(\tau, q), \qquad \tau^* = \operatorname*{argmax}_{\tau \in T} \min_{q_X \in Q_X} F(\tau, q_X).$$

By definition (1), the risk is a linear function of the probabilities $q_X(y' \,|\, x)$, which form a strategy q_X. Consequently, the function F is also a linear function of q_X for any fixed weight function τ. With a fixed strategy q_X, the function F is a linear function of weights $\tau(\theta)$. The set T of all weight functions and the set Q_X of all strategies are closed and convex. Due to the well-known duality theorem [2, 3, 10], for such functions F we have

$$\max_{\tau \in T} \min_{q_X \in Q_X} F(\tau, q_X) = F(\tau^*, q_X^*) = \min_{q_X \in Q_X} \max_{\tau \in T} F(\tau, q_X).$$

It is obvious that the equality $F(\tau, q_X^0) = 0$ holds for every $\tau \in T$. Therefore the inequality $\min\limits_{q_X \in Q_X} F(\tau, q_X) \le 0$ holds for every $\tau \in T$ and, consequently,

$$\max_{\tau \in T} \min_{q_X \in Q_X} F(\tau, q_X) = F(\tau^*, q_X^*) \le 0.$$

Since the value $F(\tau^*, q_X^*)$ is not positive, only two cases are possible: either $F(\tau^*, q_X^*) < 0$ or $F(\tau^*, q_X^*) = 0$. Let us consider the case $F(\tau^*, q_X^*) < 0$. We have the following chain:

$$\begin{aligned} F(\tau^*, q_X^*) &= \min_{q_X \in Q_X} \max_{\tau \in T} F(\tau, q_X) = \max_{\tau \in T} F(\tau, q_X^*) = \\ &= \max_{\tau \in T} \sum_{\theta \in \Theta} \tau(\theta)[R_X(q_X^*, \theta) - R_X(q_X^0, \theta)] = \max_{\theta \in \Theta} [R_X(q_X^*, \theta) - R_X(q_X^0, \theta)]. \end{aligned}$$

The chain results in the inequality

$$\max_{\theta \in \Theta} [R_X(q_X^*, \theta) - R_X(q_X^0, \theta)] < 0.$$

Consequently, the inequality $R_X(q_X^*, \theta) < R_X(q_X^0, \theta)$ holds for every model $\theta \in \Theta$ and property (3) is proved. Note that property (2) does not hold in this case.

Let us consider the case $F(\tau^*, q_X^*) = 0$. We have the following chain:

$$\begin{aligned} F(\tau^*, q_X^*) &= \max_{\tau \in T} \min_{q_X \in Q_X} F(\tau, q_X) = \min_{q_X \in Q_X} F(\tau^*, q_X) = \\ &= \min_{q_X \in Q_X} \sum_{\theta \in \Theta} \tau^*(\theta)[R_X(q_X, \theta) - R_X(q_X^0, \theta)] = \\ &= \min_{q_X \in Q_X} \left[\sum_{\theta \in \Theta} \tau^*(\theta) R_X(q_X, \theta) \right] - \sum_{\theta \in \Theta} \tau^*(\theta) R_X(q_X^0, \theta). \end{aligned}$$

It implies the equality

$$\min_{q_X \in Q_X} \sum_{\theta \in \Theta} \tau^*(\theta) R_X(q_X, \theta) = \sum_{\theta \in \Theta} \tau^*(\theta) R_X(q_X^0, \theta)$$

and thus property (2) is proved. Note that inequality (3) does not hold in this case. ■

The theorem states that any reasonable strategy for complex object recognition has to minimize the weighted sum of risks and to be of the form

$$q^* = \operatorname*{argmin}_{q \in Q} \sum_{\theta \in \Theta} \tau(\theta) R(q, \theta)$$

for some weights $\tau(\theta)$. Therefore, it has to make the decision of the form

$$y^* = \operatorname*{argmin}_{y' \in Y} \sum_{y \in Y} \Big[\sum_{\theta \in \Theta} \tau(\theta) p(x, y; \theta) \Big] w(y, y') \tag{4}$$

for a signal x. For the special case of the loss function given by

$$w(y, y') = \begin{cases} 0 & \text{if } y = y' \\ 1 & \text{if } y \neq y' \end{cases}, \tag{5}$$

decision (4) simplifies to

$$y^* = \operatorname*{argmax}_{y \in Y} \sum_{\theta \in \Theta} \tau(\theta)\, p(x, y; \theta). \tag{6}$$

The Theorem 1 does not assign weight function τ that has to be used in (4) or (6) and so it does not uniquely define a strategy. Selecting an appropriate strategy depends on requirements that have to be satisfied in such or other applied situation. The theorem outlines only a strategy subset where one has to look for an appropriate strategy. However, the theorem results in important negative conclusion. It shows which strategies must not be used in any application. For some complex objects, commonly used maximum likelihood strategy that makes a decision

$$y^* = \operatorname*{argmax}_{y \in Y} \max_{\theta \in \Theta} p(x, y; \theta) \tag{7}$$

occurs to be just such improper strategy. There are quite ordinary and realistic cases when it is impossible to express the strategy (7) in a form (6).

Example 2. Let us consider a complex object with a set $\Theta = \{1, 2\}$ of models, a set $Y = \{A, B\}$ of states and a set $X = \mathbb{R}^2$ of signals. A priori probabilities $p_Y(y = A)$ and $p_Y(y = B)$ of states are equal and do not depend on the model. For the fixed state and model, the signal $x = (x^{(1)}, x^{(2)})$ is a two-dimensional Gaussian random variable with independent components. The variances of both components $x^{(1)}$, $x^{(2)}$ are $\sigma^2 = 1$ and depend neither on the state nor on the

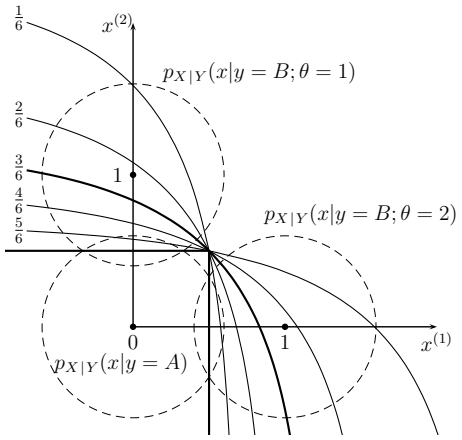

Figure 1: Family of different Bayesian strategies.

model. Expected value $\mu(y,\theta)$ of the signal $x=(x^{(1)},x^{(2)})$ depends both on the state and on the model so that $\mu(y=A,\theta=1)=\mu(y=A,\theta=2)=(0,0)$, $\mu(y=B,\theta=1)=(0,1)$, $\mu(y=B,\theta=2)=(1,0)$. Probability density distributions $p_{X\,|\,Y}(x|y;\theta)$ are shown on Figure 1 with circles.

If a loss function is given by (5) the risk $R_X(q_X,\theta)$ is a probability of a wrong decision made by strategy q_X provided θ is a true model. Let us assume such loss function and compare the maximum likelihood strategy (7) with a family (6) of Bayesian strategies. Each strategy in the family (6) is specified by concrete values of weights $\tau(1)$ and $\tau(2)$.

Maximum likelihood strategy (7) makes the decision $y^*=A$ if the observed signal $x=(x^{(1)},x^{(2)})$ belongs to the subset $X_A=\left\{(x^{(1)},x^{(2)})\in X|x^{(1)}\le 0.5,x^{(2)}\le 0.5\right\}$ and makes the decision $y^*=B$ for signals from $X_B=X\setminus X_A$. The strategy is represented at Figure 1 with a border between X_A and X_B. The border consists of two half-lines, horizontal and vertical. In a similar way, several Bayesian strategies from the family (6) are represented with curves. The strategies are built for values $\frac{1}{6},\frac{2}{6},\frac{3}{6},\frac{4}{6},\frac{5}{6}$ of the weight $\tau(1)$ and corresponding values of the weight $\tau(2)=1-\tau(1)$.

One can see that maximum likelihood strategy for recognition of the considered complex object differs from any Bayesian strategy and due to Theorem 1 is improper. In turn, it means that some Bayesian strategy exists that strongly predominates maximum likelihood strategy for both models $\theta\in\Theta$. The Theorem 1 does not say anything on how this better strategy should be selected from the family (6). However, one can guess that it is a Bayesian strategy with $\tau(1)=\tau(2)=0.5$. Really, if a true model is $\theta=1$ the maximum likelihood strategy makes a wrong decision with probability ≈ 0.37. The same is the probability of a wrong decision if $\theta=2$. Bayesian strategy with $\tau(1)=\tau(2)=0.5$ makes a wrong decision with probability ≈ 0.35 both for $\theta=1$ and for $\theta=2$. ■

We have excluded from consideration all obviously bad strategies, which left us only with Bayesian strategies. This does not mean that every Bayesian strategy predominates any non-Bayesian one. It only means that for each non-Bayesian strategy some better strategy exists and this better strategy is Bayesian. Moreover, it does not mean that each Bayesian strategy in itself is suitable for any application. Selecting a particular strategy from the set of all Bayesian strategies depends on the requirements on the strategy.

Perhaps, the most popular requirement in the theory and practice of complex hypothesis discrimination is the following so-called minimax requirement [13]. For given sets X, Y, Θ and probabilities $p(x, y; \theta)$, $x \in X$, $y \in Y$, $\theta \in \Theta$, a strategy $q_X \colon Y \times X \to \mathbb{R}$ has to be found that minimizes c subject to the condition that $R_X(q_X, \theta) \leq c$ for each $\theta \in \Theta$. In other words, the strategy

$$q_X^* = \operatorname*{argmin}_{q_X \in Q_X} \max_{\theta \in \Theta} R_X(q_X, \theta) \tag{8}$$

is called a minimax strategy. Obviously, the minimax strategy is not improper and, consequently, is Bayesian. Indeed, if q_X^* would be improper then such strategy q'_X would exist that the inequalities

$$R_X(q'_X, \theta) < R_X(q_X^*, \theta), \quad \theta \in \Theta,$$

would be valid as well as the inequality

$$\max_{\theta \in \Theta} R_X(q'_X, \theta) < \max_{\theta \in \Theta} R_X(q_X^*, \theta)$$

that would contradict (8). Minimax strategies are fruitfully used in the theory and practice of recognition without learning [7, 22].

However, the minimax requirement is not the only possible meaningful one. In multi-criteria decision making [23] another reasonable requirement is used that differs from the minimax one. Let us follow the way shown in [23] and define a number $\min\limits_{q_X \in Q_X} R_X(q_X, \theta)$ for each $\theta \in \Theta$ and call it the risk of optimal strategy. In general, the risk of the optimal strategy depends on the model and so does the optimal strategy $\operatorname*{argmin}\limits_{q_X \in Q_X} R_X(q_X, \theta)$. However, for some extraordinary convenient complex objects the strategy $\operatorname*{argmin}\limits_{q_X \in Q_X} R_X(q_X, \theta)$ does not depend on the model.

Definition 5. *A strategy q_X^* satisfying the equality $R_X(q_X^*, \theta) = \min\limits_{q_X \in Q_X} R_X(q_X, \theta)$ for every model $\theta \in \Theta$ is called an overall optimal strategy.*

The existence of an overall optimal strategy is a rare exception rather than the rule. We formulate the problem of complex object recognition so that its solution is an overall optimal strategy when such strategy exists. Otherwise, the obtained strategy risk deviates from the optimal strategy risk as little as possible. More exactly, for a given complex object $\langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle$

and a loss function w: $Y \times Y \to \mathbb{R}$, the strategy q_X: $Y \times X \to \mathbb{R}$ has to be found that minimizes c subject to conditions

$$R_X(q_X, \theta) - \min_{q'_X \in Q} R_X(q'_X, \theta) \le c, \quad \theta \in \Theta.$$

Definition 6. *The strategy* $q^*_X = \operatorname*{argmin}_{q_X \in Q_X}\big[\max_{\theta \in \Theta}\big(R_X(q_X, \theta) - \min_{q'_X \in Q} R_X(q'_X, \theta)\big)\big]$ *is called a closest to optimal strategy.*

Just as the minimax strategy the closest to optimal strategy is not improper and, consequently, is also Bayesian. The concept of the closest to optimal strategy is a straightforward application of the recommendations in [23] to our special case. On the other hand, the closest to optimal strategy is a generalization of the minimax deviation (regret) strategy from [1]. The latter is a special case of the closest to optimal strategy when the statistical model is defined up to a priori probabilities of states.

The minimax approach (8) and the closest to optimal one (6) result in different Bayesian strategies. If the task is only recognition without learning, there is no reason to decide which one of these two approaches is better. However, the situation is different when the approaches are applied to recognition with learning. In this case the minimax strategies are not suitable for recognition with learning at all. They simply ignore the learning sample. This defect of minimax approach has been detected by H.Robbins at the very beginning of his empirical Bayesian approach. This defect is considered in details in the next Section 3. Then in the Section 4, we show that the closest to optimal strategies are free of this defect.

3 Empirical Bayesian approach and unsupervised learning

The principal property of a complex object is that its model is unknown but it is not random. It is fixed and does not change. The property gives rise to conjecture that recognizing a sample of independent and identically distributed objects can be performed better than isolated recognition of each element in the sample independently. The idea was researched by H.Robbins [16] decades ago and initiated an empirical Bayesian approach [15]. In fact, the same idea is the basis of an approach in pattern recognition known as unsupervised learning. However, these two approaches exist as if they arouse and were being developed independently of one another. We consider them as two different paths to the same goal and formulate a problem that may be considered as one of possible concretizations of empirical Bayesian approach and one of possible modifications of unsupervised learning.

H.Robbins explains his idea on the following simple example [16], which we will use several times through the article. An object is considered with a set $X = \mathbb{R}$ of signals, a set $Y = \{1, 2\}$ of states and a set $\Theta = \{\theta \mid 0 \le \theta \le 1\}$ of

models. The object generates a Gaussian random signal $x \in X$, whose variance is 1, and whose mean equals 1 if the object is in state 1 and equals (-1) if the object is in state 2. Only a priori probabilities of the states are unknown. They depend on the model so that the probability of state 1 is θ and the probability of state 2 is $1-\theta$. In other words, the probability densities $p(x, y; \theta)$ that define the complex object are

$$p(x, y; \theta) = \frac{\theta_y}{\sqrt{2\pi}} \mathrm{e}^{-\frac{1}{2}(x-\mu_y)^2}, \tag{9}$$

where $\theta_1 = \theta$, $\theta_2 = 1-\theta$, $\mu_1 = 1$, $\mu_2 = (-1)$. The loss function is given by (5). Let us consider optimal and minimax strategies

$$q_X^{\mathrm{opt}}(\theta) = \operatorname*{argmin}_{q_X \in Q_X} R_X(q_X, \theta), \qquad q_X^{\mathrm{minmax}} = \operatorname*{argmin}_{q_X \in Q_X} \max_{\theta \in \Theta} R_X(q_X, \theta)$$

and see how the risks $R_X(q_X^{\mathrm{opt}}(\theta), \theta)$ and $R_X(q_X^{\mathrm{minmax}}, \theta)$ depend on the model θ (see Figure 2). Each strategy has a drawback compared to the other one. The strategy $q_X^{\mathrm{opt}}(\theta)$ can be used only if the model θ is known. The strategy q_X^{minmax} can be used even if the model θ is unknown, however the risk $R_X(q_X^{\mathrm{minmax}}, \theta)$ is worse than $R_X(q_X^{\mathrm{opt}}(\theta), \theta)$ if θ is near to 1 or 0.

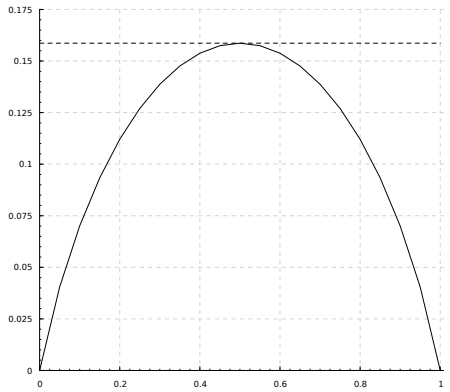

Figure 2: The dependency of $R_X(q_X^{\mathrm{minmax}}, \theta)$ on model θ is plotted with a dashed curve. Probabilities $\min_{q_X \in Q_X} R_X(q_X, \theta)$ are plotted with a solid curve.

Suppose that rather than a single object, a sample of independent random objects has to be recognized. The model θ according to which the objects are generated is unknown but fixed for the whole sample. H.Robbins formulates the main idea of empirical Bayesian approach that such a sample can be recognized much better than each of its elements independently. His important negative result is that this improvement cannot be achieved in the framework of the minimax approach. Let us state this negative result exactly.

Let a complex object be defined by the sets $X = \mathbb{R}$, $Y = \{1, 2\}$, $\Theta = \{\theta \mid 0 \le \theta \le 1\}$ and the probability density (9). Let the loss function have the form (5). For this complex object and a positive integer n, another complex object is defined, a so called compound complex object. It is characterized by sets X^n, Y^n, Θ, probabilities

$$p^n(\bar{x}, \bar{y}; \theta) = \prod_{i=1}^{n} p(x_i, y_i; \theta), \quad \bar{x} \in X^n,\ \bar{y} \in Y^n,\ \theta \in \Theta,$$

and losses

$$w^n(\bar{y}, \bar{y'}) = \frac{1}{n} \sum_{i=1}^{n} w(y_i, y'_i), \quad \bar{y} \in Y^n,\ \bar{y'} \in Y^n.$$

Let $q_{X^n}\colon Y^n \times X^n \to \mathbb{R}$ be the recognition strategy for the compound object and let

$$R_{X^n}(q_{X^n}, \theta) = \sum_{\bar{y'} \in Y^n} \sum_{\bar{y} \in Y^n} \sum_{\bar{x} \in X^n} q_{X^n}(\bar{y'} \,|\, \bar{x})\, p^n(\bar{x}, \bar{y}; \theta)\, w^n(\bar{y}, \bar{y'})$$

be the risk of the strategy q_{X^n} with respect to the model θ. The quoted negative result states that

$$\min_{q_{X^n}} \max_{\theta \in \Theta} R_{X^n}(q_{X^n}, \theta) = \min_{q_X} \max_{\theta \in \Theta} R_X(q_X, \theta)$$

for any number n. It means that the quality of the strategy

$$q_{X^N}^{\text{minmax}} = \operatorname*{argmin}_{q_{X^n}} \max_{\theta \in \Theta} R_{X^n}(q_{X^n}, \theta)$$

is not better than

$$q_X^{\text{minmax}} = \operatorname*{argmin}_{q_X} \max_{\theta \in \Theta} R_X(q_X, \theta).$$

Minimax strategies do not take advantage of the situation that multiple objects with a common model have to be recognized. We will show in Section 4 that this drawback of minimax strategies also holds in a more general case.

Despite of this clearly negative result, H.Robbins presents a heuristic strategy that for the objects of the form (9) and observed signals $(x_1, x_2, \ldots, x_n)$ makes the decisions

$$y_i = \begin{cases} 1 & \text{if } x_i \ge \alpha \\ 2 & \text{if } x_i < \alpha \end{cases} \quad \text{where} \quad \alpha = \frac{1}{2} \ln \frac{n - \sum_{i=1}^{n} x_i}{n + \sum_{i=1}^{n} x_i}. \tag{10}$$

The strategy (10) does not predominate $q_X^{\text{opt}}(\theta)$ because $q_X^{\text{opt}}(\theta)$ is optimal by definition. The strategy (10) does not predominate $q_{X^n}^{\text{minmax}}$ because $q_{X^n}^{\text{minmax}}$ is Bayesian and no strategy predominates it. However, the strategy (10) can be used even when the model θ is unknown. From this point of view it is better than the strategy $q_X^{\text{opt}}(\theta)$ that makes the decision

$$y_i = \begin{cases} 1 & \text{if } x_i \ge \alpha \\ 2 & \text{if } x_i < \alpha \end{cases} \quad \text{where} \quad \alpha = \frac{1}{2} \ln \frac{1 - \theta}{\theta} \tag{11}$$

that depends on θ. The strategy (10) is better than $q_{X^n}^{\text{minmax}}$ in a sense that its quality improves as n increases while the quality of $q_{X^n}^{\text{minmax}}$ does not depend on n and remains equal to the quality of q_X^{minmax}. Moreover, for large enough n the strategy (10) behaves almost like $q_X^{\text{opt}}(\theta)$.

The strategy (10) illustrates the main idea of the empirical Bayesian approach [15, 16, 17], namely that an object sample can be recognized much better than by recognizing each object independently. However, the example says nothing about the requirement that the strategy (10) should be derived from. It only makes clear that it is not the minimax one. The problem of compound object recognition is not stated as precisely as the similar problems in the Bayesian or minimax decision theory. Therefore it is unclear how the approach should be formulated in general, not only for the special case (9).

One of the possible specifications of the empirical Bayesian approach may be the following problem definition of compound complex object recognition. The input data of this problem consist of a quadruple $\langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle$ representing a complex object, a loss function w: $Y \times Y \to \mathbb{R}$ and a positive integer n. These data define risks

$$\min_{q_X \in Q_X} R_X(q_X, \theta) = \min_{q_X \in Q_X} \sum_{y' \in Y} \sum_{y \in Y} \sum_{x \in X} q_X(y' \,|\, x)\, p(x, y; \theta)\, w(y, y')$$

of optimal strategies, a compound object $\langle X^n, Y^n, \Theta, p^n : X^n \times Y^n \times \Theta \to \mathbb{R} \rangle$ and a loss function w^n: $Y^n \times Y^n \to \mathbb{R}$ where

$$p^n(\bar{x}, \bar{y}; \theta) = \prod_{i=1}^{n} p(x_i, y_i; \theta), \quad \bar{x} \in X^n,\ \bar{y} \in Y^n,\ \theta \in \Theta,$$

$$w^n(\bar{y}, \bar{y'}) = \frac{1}{n} \sum_{i=1}^{n} w(y_i, y'_i), \quad \bar{y} \in Y^n,\ \bar{y'} \in Y^n.$$

Let q_{X^n}: $Y^n \times X^n \to \mathbb{R}$ be a randomized strategy that makes a decision about hidden states $y_1, y_2, \cdots, y_n$ on the base of observed signals $y_1, y_2, \cdots, y_n$. Let Q_{X^n} be the set of all such strategies. For $q_{X^n} \in Q_{X^n}$ and model $\theta \in \Theta$, the risk is a number

$$R_{X^n}(q_{X^n}, \theta) = \sum_{\bar{y'} \in Y^n} \sum_{\bar{y} \in Y^n} \sum_{\bar{x} \in X^n} q_{X^n}(\bar{y'} \,|\, \bar{x})\, p^n(\bar{x}, \bar{y}; \theta)\, w^n(\bar{y}, \bar{y'}).$$

The problem of compound complex object recognition is defined so that for a quadruple

$$\langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle,$$

a loss function w: $Y \times Y \to \mathbb{R}$ and a positive integer n the closest to optimal strategy

$$q_{X^n}^* = \operatorname*{argmin}_{q_{X^n} \in Q_{X^n}} \max_{\theta \in \Theta} \big[R_{X^n}(q_{X^n}, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta) \big] \tag{12}$$

has to be found. Obviously, the solution $q^*_{X^n}$ to this problem is a Bayesian strategy.

Widely known unsupervised learning is another approach that copes with recognizing compound complex objects. It proceeds as follows. First, for given sets X, Y, Θ, probabilities $p(x, y; \theta)$, $x \in X$, $y \in Y$, $\theta \in \Theta$, and a sample $\bar{x} = (x_1, x_2, \ldots, x_n)$ the maximum likelihood estimate

$$\theta^{\mathrm{ML}} = \operatorname*{argmax}_{\theta \in \Theta} \prod_{i=1}^{n} \sum_{y \in Y} p(x_i, y; \theta) = \operatorname*{argmax}_{\theta \in \Theta} \sum_{i=1}^{n} \log \sum_{y \in Y} p(x_i, y; \theta) \tag{13}$$

is found. Then, the elements of the sample are recognized using the strategy

$$q^{\mathrm{ML}}_X = \operatorname*{argmin}_{q_X \in Q_X} R_X(q_X, \theta^{\mathrm{ML}}) \tag{14}$$

as if the obtained model θ^{ML} would be a true model.

The maximum likelihood estimation (13) is of the main interest here because it is not trivial even in the simplest cases. To solve (13), the algorithms [20] are used, which later became widely known as the EM-algorithms [6]. The problem (13) is significant in its own right. Nevertheless, for some complex objects the strategy (13, 14) is improper even though each strategy of the family

$$q^*_X = \operatorname*{argmin}_{q_X \in Q_X} R_X(q_X, \theta^*), \quad \theta^* \in \Theta, \tag{15}$$

is Bayesian. Each strategy of the family (15) differs from the strategy (13), (14) at least due to the fact that their formats differ from one another, the first being $Y \times X \to \mathbb{R}$ and the second $Y^n \times X^n \to \mathbb{R}$. However, even if $n = 1$ and their formats become identical, some complex objects exist such that the strategy (13), (14) differs from each strategy of the family (15). Moreover, it differs from every Bayesian strategy.

Example 3. Let us consider a complex object and a loss function as in the Example 2 and compare the unsupervised learning strategy (13), (14) with strategies of the family (15) for the simplest case when $n = 1$. For a given signal $x = (x^{(1)}, x^{(2)})$ the strategy (15) makes a decision

$$y^* = \operatorname*{argmax}_{y \in Y} p(x^{(1)}, x^{(2)}, y; \theta^*), \tag{16}$$

that depends on θ^*. For $\theta^* = 1$ the strategy q^*_X divides the plane $X = \mathbb{R}^2$ into two half-planes $X^*(A, \theta^* = 1)$ and $X^*(B, \theta^* = 1)$ with a horizontal line with coordinate $x^{(2)} = 0.5$. For $\theta^* = 2$ it divides X into $X^*(A, \theta^* = 2)$ and $X^*(B, \theta^* = 2)$ with a vertical line with coordinate $x^{(1)} = 0.5$ (see Figure 3). The strategy (13), (14) makes a decision

$$y^{\mathrm{ML}} = \operatorname*{argmax}_{y \in Y} p(x^{(1)}, x^{(2)}, y; \operatorname*{argmax}_{\theta \in \Theta} \sum_{y \in Y} p(x^{(1)}, x^{(2)}, y; \theta)) \tag{17}$$

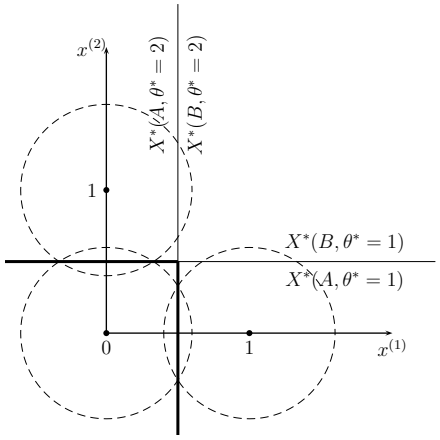

Figure 3: Maximum likelihood strategy.

that does not depend on θ. This decision is A on subset

$$\{(x^{(1)},x^{(2)})\in X | x^{(1)}\le 0.5, x^{(2)}\le 0.5\} = X^*(A,\theta^*=1)\cap X^*(A,\theta^*=2)$$

that differs both from $X^*(A,\theta^*=1)$ and from $X^*(A,\theta^*=2)$. Therefore, the strategy (13), (14) differs both from $\underset{q_X\in Q_X}{\operatorname{argmin}} R_X(q_X,\theta^*=1)$ and from $\underset{q_X\in Q_X}{\operatorname{argmin}} R_X(q_X,\theta^*=2)$. It has already been shown in Example 2 that the strategy differs from any Bayesian strategy. ■

The heuristic strategy (10) by H.Robbins, the closest to optimal strategy (12) and the unsupervised learning strategy (13, 14) are three different strategies. They are compared experimentally for the complex model (9) in Example 9 in Section 6. In this experiment, *the closest to optimal strategy yields the best result, the unsupervised learning strategy the worst, and the heuristic strategy an intermediate one.*

4 Recognition with learning

As in the previous sections, we consider complex objects, defined by a quadruple $\langle X,Y,\Theta,p:X\times Y\times\Theta\to\mathbb{R}\rangle$. In addition, a certain source of data is considered. The source generates the so called learning information z that belongs to a finite set Z. The learning information z is random and depends on the model $\theta\in\Theta$, so that the probability $p(z;\theta)$ is defined for every $z\in Z$ and $\theta\in\Theta$. It is crucial that for a fixed model θ, the learning information z depends neither on the current state y nor on the current signal x of the object, so that $p(z,x,y;\theta)=p(z;\theta)p(x,y;\theta)$. Formally, the source of learning information is defined with a triple $\langle Z,\Theta,p:Z\times\Theta\to\mathbb{R}\rangle$.

Usually, learning information is obtained in a following way. The object is put into a special condition so that even its state becomes observable. Under this condition, a sample $(x_1, y_1; x_2, y_2; \dots\ x_n, y_n)$ of n independent random pairs (x, y) is observed so that $Z = (X \times Y)^n$ and

$$p(z;\theta) = p(x_1, y_1, x_2, y_2, \dots, x_n, y_n; \theta) = \prod_{i=1}^{n} p(x_i, y_i; \theta).$$

Our formulation is more general. In case of unsupervised learning, the learning information z is an element of X^n. When the learning information is absent, we introduce a special element 0 and define $Z = \{0\}$ and $p(z = 0; \theta) = 1$ for every model $\theta \in \Theta$. The learning information can have a different nature, for example, it can be an expert opinion about the model. Moreover, the learning information can come from several independent sources simultaneously as it is mentioned in [12]. In this case, the set Z_s and the probabilities $p(z; \theta, s)$, $z \in Z_s$, $\theta \in \Theta$, $s \in S$, are given for every source $s \in S$.

The problem of recognition with learning means that for a given complex object and a learning information source a strategy has to be developed that makes a decision about current object state based both on the learning information and on the current signal generated by the object. We express this statement in the form

$$\begin{gathered} \langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle, \langle Z, \Theta, p : Z \times \Theta \to \mathbb{R} \rangle \\ \downarrow \\ q_{X \times Z} : Y \times X \times Z \to \mathbb{R}. \end{gathered} \tag{18}$$

The value $q_{X \times Z}(y'|x, z)$ of the strategy $q_{X \times Z}$ is a conditional probability of the decision "a state of an object is y'" under condition that a signal x is observed and learning information z is available. The form

$$\begin{gathered} \langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle \\ \downarrow \\ q_X : Y \times X \to \mathbb{R}. \end{gathered} \tag{19}$$

expresses the problem of recognition without learning. One can see that the form (18) is a special case of (19). Really, (18) can be equivalently expressed as

$$\begin{gathered} \langle X^*, Y, \Theta, p^* : X^* \times Y \times \Theta \to \mathbb{R} \rangle \\ \downarrow \\ q_{X^*} : Y \times X^* \to \mathbb{R} \end{gathered}$$

with $X^* = X \times Z$ and $p^*(x^*, y; \theta) = p^*(x, z, y; \theta) = p(z; \theta) \cdot p(x, y; \theta)$. In other words, a complex object with a source of learning information can be replaced with another object that generates specific signals composed of two components. The first component depends both on the current state of the object and on the

model. The second one depends directly only on the model and under a fixed model depends neither on the current state nor on the first component.

Theorem 1, proved in Section 2 in the general form, remains valid for recognition with learning. However, for subsequent considerations another formulation of the theorem is more convenient. A strategy $q_{X\times Z} : Y \times X \times Z \to \mathbb{R}$ will be represented as a superposition of two functions: a learning procedure $g_Z\colon Z \to Q$ and a recognition strategy $q_X\colon Y \times X \to \mathbb{R}$. Given learning information z, the learning procedure g_Z defines the recognition strategy $g_Z(z)\colon Y \times X \to R$. This means that after learning with random information z and observing a signal x, the decision "the object is in state y'" is made with probability $g_Z(z)(y'\,|\,x)$. Let G_Z be the set of all possible learning procedures of the form $Z \to Q$.

As before, $R_X(q_X, \theta)$ is the risk of recognition strategy $q_X \in Q_X$ with respect to the model $\theta \in \Theta$. The number $R_X(g_Z(z), \theta)$ is the risk of the strategy $g_Z(z)$, obtained by applying the learning procedure $g_Z\colon Z \to Q$ to the learning information $z \in Z$. The risk $R_X(g_Z(z), \theta)$ is random because it depends on the random information z. Let $R_Z(g_Z, \theta)$ denote the expectation of the risk over the set of all learning informations,

$$R_Z(g_Z, \theta) = \sum_{z\in Z} p(z;\theta) R_X(g_Z(z), \theta).$$

Definition 7. *A learning procedure g_Z^* is called Bayesian if a weight function $\tau\colon \Theta \to \mathbb{R}$ exists such that*

$$g_Z^* = \operatorname*{argmin}_{g_Z\in G_Z} \sum_{\theta\in\Theta} \tau(\theta) R_Z(g_Z, \theta).$$

Now the restriction of Theorem 1 to learning procedures reads as follows.

Theorem 2. *For every learning procedure $g_Z^0 \in G$, either a weight function $\tau^* \in T$ exists such that*

$$g_Z^0 = \operatorname*{argmin}_{g_Z\in G_Z} \sum_{\theta\in\Theta} \tau^*(\theta) R_Z(g_Z, \theta)$$

or a learning procedure g_Z^* exists such that for every model $\theta \in \Theta$ the inequality

$$R_Z(g_Z^*, \theta) < R_Z(g_Z^0, \theta)$$

holds. These two properties of learning procedures are incompatible.

The theorem says that every learning procedure is either Bayesian or improper. In the special case when a sample $(x_1, y_1; x_2, y_2; \ldots; x_n, y_n)$ is the learning information and a signal x_0 is observed, the decision y^* about the current object state has to be of the form

$$y^* = \operatorname*{argmin}_{y'\in Y} \sum_{y_0\in Y} \sum_{\theta\in\Theta} \tau(\theta) \prod_{i=0}^{n} p(x_i, y_i; \theta) w(y_0, y') \tag{20}$$

with some weights $\tau(\theta)$, $\theta \in \Theta$. For the loss function (5), the strategy is simplified to

$$y^* = \underset{y_0 \in Y}{\operatorname{argmax}} \sum_{\theta \in \Theta} \tau(\theta) \prod_{i=0}^{n} p(x_i, y_i; \theta).$$

Strategies that are commonly used in pattern recognition have other form. The most widely known is the method based on maximum likelihood estimation [7, 22]. According to this method, the model

$$\theta^{\mathrm{ML}} = \underset{\theta \in \Theta}{\operatorname{argmax}} \prod_{i=1}^{n} p(x_i, y_i; \theta) \tag{21}$$

is found and then for an observed signal x the decision

$$y^{\mathrm{ML}} = \underset{y' \in Y}{\operatorname{argmin}} \sum_{y \in Y} p(x, y; \theta^{\mathrm{ML}}) w(y, y') \tag{22}$$

is made, as if the maximum likelihood model was identical to the true one. For some complex objects, the strategy (21, 22) differs from (20) and thus it is improper. One of such complex objects was considered in Examples 2 and 3. The following example shows more sharply the main defect of maximum likelihood learning (21), (22).

Example 4. Let us return to complex object $\langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle$ described in a Section 3 with sets $X = \mathbb{R}$, $Y = \{1, 2\}$, $\Theta = \{\theta | 0 \leq \theta \leq 1\}$ and probability density distribution $p : X \times Y \times \Theta \to R$ of the form (9). The loss function is given by (5). If no learning sample is available one can use a minimax recognition strategy that does not depend on a model. The strategy decides that $y = 1$ if $x \geq 0$ and $y = 2$ if $x < 0$. The strategy makes a wrong decision with probability ≈ 0.16 independently of a model and seems to be quite reasonable.

Suppose now that the learning sample is obtained that consists of only pair (x_0, y_0). It is very small sample but one can expect that it enables at least small improvement as compared with minimax strategy that works if no learning sample is available. However, the maximum likelihood learning (21), (22) results in a strategy that ignores a current signal x at all and decides that the object is in the state y_0 independently of signal generated by object. So, if the learning sample is of very small size it is reasonable to ignore it rather than to make use of it with maximum likelihood learning. ■

Nevertheless, the maximum likelihood learning (21, 22) has an important positive property that can be called asymptotic optimality. Asymptotic optimality qualifies a learning procedure in a similar sense as consistency qualifies a parameter estimate. We formulate and prove this property after several preliminary remarks.

First, we assume that for every two different models θ_1 and θ_2 a pair $(x, y) \in X \times Y$ exists such that $p(x, y; \theta_1) \neq p(x, y; \theta_2)$. Obviously, this assumption is not

a restriction. Really, if there are two models θ_1 and θ_2 such that $p(x,y;\theta_1) = p(x,y;\theta_2)$ for each pair $(x,y) \in X \times Y$ they can be treated as equivalent and be replaced with one model.

Then, we rely on properties of entropy-like functions $\sum\limits_{i=1}^{n} a_i \log x_i$ of $2n$ strictly positive arguments such that $\sum\limits_{i=1}^{n} a_i = \sum\limits_{i=1}^{n} x_i = 1$. However, we extrapolate such functions onto the set of non-negative numbers assuming that $a \log x = 0$ if $a = x = 0$ and $a \log x = -\infty$ if $a > 0$ and $x = 0$. Under these assumptions the difference

$$\sum_{x\in X}\sum_{y\in Y} p(x,y;\theta^*)\log p(x,y;\theta') - \sum_{x\in X}\sum_{y\in Y} p(x,y;\theta^*)\log p(x,y;\theta^*)$$

is defined for every pair (θ^*, θ') and is strictly negative if $\theta^* \neq \theta'$. The statement results from the known property of entropy-like functions that

$$\max_{x_1}\max_{x_2}\ldots\max_{x_n}\sum_{i=1}^{n} a_i \log x_i = \sum_{i=1}^{n} a_i \log a_i$$

and if $x_{i^*} \neq a_{i^*}$ for some i^* then $\sum\limits_{i=1}^{n} a_i \log a_i > \sum\limits_{i=1}^{n} a_i \log x_i$.

At last, we use a fundamental relation between the averaged value $\frac{1}{n}\sum\limits_{i=1}^{n} x_i$ and the expected value $\sum\limits_{x\in X} x \cdot p(x)$. The relation restricted to our special case has the following formulation. Let n be a positive number, X be a finite set of numbers, $p(x)$ be the probability of $x \in X$ and $Z_n = \{z \in X^n | \frac{1}{n}\sum\limits_{i=1}^{n} x_i \geq 0\}$. If $\sum\limits_{x\in X} x \cdot p(x) < 0$ then

$$\lim_{n\to\infty}\sum_{z\in Z_n}\prod_{i=1}^{n} p(x_i) = 0. \tag{23}$$

Theorem 3. *Let $Z = (X \times Y)^n$ and the procedure g_n^{ML}:$(X \times Y)^n \to Q$ be defined by*

$$g_n^{\mathrm{ML}}(z) = \operatorname*{argmin}_{q_X \in Q_X} R_X(q_X, \operatorname*{argmax}_{\theta\in\Theta} p(z;\theta)), \quad z \in (X \times Y)^n. \tag{24}$$

Then for every $\theta^* \in \Theta$

$$\lim_{n\to\infty}[R_Z(g_n^{\mathrm{ML}}, \theta^*) - \min_{q_X \in Q_X} R_X(q_X, \theta^*)] = 0. \tag{25}$$

Proof. Let us choose an arbitrary model θ^* and fix it for the whole proof. Let us denote

$$\theta^{\mathrm{ML}}(z) = \operatorname*{argmax}_{\theta\in\Theta} p(z;\theta) = \operatorname*{argmax}_{\theta\in\Theta}\prod_{i=1}^{n} p(x_i, y_i;\theta).$$

The model $\theta^{\mathrm{ML}}(z)$ is random because it depends on a random sample z. If the probability of the sample z is

$$p(z;\theta^*) = p(x_1, y_1, x_2, y_2, \dots x_n, y_n; \theta^*) = \prod_{i=1}^{n} p(x_i, y_i; \theta^*)$$

then inequality $\theta^{\mathrm{ML}}(z) \neq \theta^*$ is a random event with its probability depending on n. We prove that the probability of inequality $\theta^{\mathrm{ML}}(z) \neq \theta^*$ converges to zero when n increases. Let us define three subsets of the set $Z = (X \times Y)^n$:

$$Z(\theta^*, \theta', n) = \{z | p(z;\theta') \geq p(z;\theta^*)\}, \quad Z(\theta^*, n) = \bigcup_{\theta' \in \Theta \setminus \{\theta^*\}} Z(\theta^*, \theta', n),$$

$$Z^{\mathrm{ML}}(\theta^*, n) = \{z | \theta^{\mathrm{ML}}(z) \neq \theta^*\}$$

and denote $P(Z') = \sum\limits_{z \in Z'} p(z;\theta^*)$ for each $Z' \subset Z$. Evidently,

$$Z^{\mathrm{ML}}(\theta^*, n) \subset Z(\theta^*, n) = \bigcup_{\theta' \in \Theta \setminus \{\theta^*\}} Z(\theta^*, \theta', n)$$

that results in

$$P(Z^{\mathrm{ML}}(\theta^*, n)) \leq P(Z(\theta^*, n)) \leq \sum_{\theta' \in \Theta \setminus \{\theta^*\}} P(Z(\theta^*, \theta', n)). \tag{26}$$

The subset $Z(\theta^*, \theta', n)$ consists of all samples $z = (x_1, y_1, x_2, y_2, \dots x_n, y_n)$ that satisfy the inequality

$$\frac{1}{n} \sum_{i=1}^{n} [\log p(x_i, y_i; \theta') - \log p(x_i, y_i; \theta^*)] \geq 0.$$

The left side of the inequality is the average of n independent and identically distributed random numbers $\log p(x, y; \theta') - \log p(x, y; \theta^*)$. The mean value

$$\sum_{x \in X} \sum_{y \in Y} p(x, y; \theta^*) \log p(x, y; \theta') - \sum_{x \in X} \sum_{y \in Y} p(x, y; \theta^*) \log p(x, y; \theta^*)$$

of these numbers is negative and therefore $\lim\limits_{n \to \infty} P(Z(\theta^*, \theta', n) = 0$ due to (23) and

$\lim\limits_{n\to\infty} P(Z^{\mathrm{ML}}(\theta^*,n)) = 0$ due to (26). The proof is concluded by the chain

$$\begin{aligned}
&\lim_{n\to\infty}[R_Z(g_n^{\mathrm{ML}},\theta^*) - \min_{q_X\in Q_X} R_X(q_X,\theta^*)] = \\
&\quad = \lim_{n\to\infty} \sum_{z\in(X\times Y)^n} p(z;\theta^*)[R_X(g_n^{\mathrm{ML}}(z),\theta^*) - \min_{q_X\in Q_X} R_X(q_X,\theta^*)] \\
&\quad = \lim_{n\to\infty} \sum_{\substack{z\in(X\times Y)^n \\ \theta^{\mathrm{ML}}(z)\neq\theta^*}} p(z;\theta^*)[R_X(g_n^{\mathrm{ML}}(z),\theta^*) - \min_{q_X\in Q_X} R_X(q_X,\theta^*)] \\
&\quad \le \lim_{n\to\infty} \sum_{\substack{z\in(X\times Y)^n \\ \theta^{\mathrm{ML}}(z)\neq\theta^*}} p(z;\theta^*)[\max_{y\in Y}\max_{y'\in Y} w(y,y') - \min_{y\in Y}\min_{y'\in Y} w(y,y')] \\
&\quad = \lim_{n\to\infty} [\max_{y\in Y}\max_{y'\in Y} w(y,y') - \min_{y\in Y}\min_{y'\in Y} w(y,y')] \sum_{\substack{z\in(X\times Y)^n \\ \theta^{\mathrm{ML}}(z)\neq\theta^*}} p(z;\theta^*) \\
&\quad = [\max_{y\in Y}\max_{y'\in Y} w(y,y') - \min_{y\in Y}\min_{y'\in Y} w(y,y')] \lim_{n\to\infty} P(Z_n^{\mathrm{ML}}(\theta^*,n)) = 0.
\end{aligned}$$

■

We restrict our further considerations to Bayesian learning procedures, focusing in particular on the minimax procedure $\operatorname{argmin}\limits_{g_Z\in G_Z}\max\limits_{\theta\in\Theta} R_Z(g_Z,\theta)$ and the closest to optimal procedure $\operatorname{argmin}\limits_{g_Z\in G_Z}\max\limits_{\theta\in\Theta}\big[R_Z(g_Z,\theta) - \min\limits_{q_X\in Q_X} R_X(q_X,\theta)\big]$. Both procedures belong to the Bayesian class. However, minimax procedures have a fundamental drawback that they sometimes do not make use of a learning sample, no matter how large it is. For a rather wide class of complex objects, the minimax learning procedure simply ignores the learning sample.

Theorem 4. *Let for a complex object $\langle X,Y,\Theta,p: X\times Y\times\Theta\to\mathbb{R}\rangle$ such a model θ^* and a strategy q_X^* exist that*

$$q_X^* = \operatorname*{argmin}_{q_X\in Q_X} R_X(q_X,\theta^*), \quad \theta^* = \operatorname*{argmax}_{\theta\in\Theta} R_X(q_X^*,\theta).$$

Then any source $\langle Z,\Theta,p: Z\times\Theta\to\mathbb{R}\rangle$ of learning information and any learning procedure $g_Z\colon Z\to Q$ satisfy the inequality

$$\max_{\theta\in\Theta} R_Z(g_Z,\theta) \ge \max_{\theta\in\Theta} R_X(q_X^*,\theta). \tag{27}$$

Proof. For any learning procedure g_Z, we have the chain

$$\begin{aligned}
&\max_{\theta\in\Theta} R_Z(g_Z,\theta) \ge R_Z(g_Z,\theta^*) = \sum_{z\in Z} p(z;\theta^*) R_X(g_Z(z),\theta^*) \ge \\
&\ge \sum_{z\in Z} p(z;\theta^*) R_X(q_X^*,\theta^*) = R_X(q_X^*,\theta^*) = \max_{\theta\in\Theta} R_X(q_X^*,\theta).
\end{aligned}$$

■

The theorem shows that there are complex objects for which the minimax approach is particularly inappropriate. Inequality (27) states that any learning procedure, however sophisticated it may be, is useless from the minimax point of view. It cannot yield a recognition strategy that would be better than some strategy that does not use the learning sample at all. The following examples show that such situations are nothing unusual.

Example 5. The conditions of Theorem 4 are satisfied if for every weight function $\tau \in T$, the model set Θ contains a model θ' such that $p(x, y; \theta') = \sum\limits_{\theta \in \Theta} \tau(\theta) p(x, y; \theta)$. One of such model sets is just the set (9), for which H.Robbins shows that the minimax strategy does not achieve the result that the empirical Bayesian approach does. ■

Example 6. Let $X = \mathbb{R}^2$, $Y = \{1, 2\}$, $p_Y(y = 1) = p_Y(y = 2) = 0.5$, and $p_{X|Y}(x|y)$ be a two-dimensional Gaussian probability density distribution with the unit covariance matrix. The expected values θ_y, $y \in \{1, 2\}$, of the random signal $x \in X$ are unknown. It is only known that they belong to given closed convex disjoint sets Θ_1 and Θ_2. This means, the set of models is $\Theta = \Theta_1 \times \Theta_2$. This model set does not satisfy the condition of Example 5 but it satisfies the conditions of Theorem 4. In this case, the model $\theta^* \in \Theta$ is the pair

$$(\theta_1^*, \theta_2^*) = \operatorname*{argmin}_{\theta_1 \in \Theta_1, \theta_2 \in \Theta_2} (\theta_1 - \theta_2)^2,$$

and the strategy q^* makes the decision

$$y^* = \begin{cases} 1 & \text{if } (x - \theta_1^*)^2 \leq (x - \theta_2^*)^2, \\ 2 & \text{if } (x - \theta_1^*)^2 > (x - \theta_2^*)^2. \end{cases}$$

It is an acceptable recognition strategy for the complex object considered. However, it is clear that a better strategy could be found if a learning sample was available, in particular if this sample was large enough. Nevertheless, the minimax strategy ignores this opportunity even if the learning sample is very large. ■

Example 7. Let p_y: $X \to \mathbb{R}$, $y \in \{1, 2\}$, be two functions whose values are conditional probabilities $p_{X|Y}(x|y)$, so that $p_y(x) = p_{X|Y}(x|y)$. It is a typical situation in pattern recognition that the functions p_1 and p_2 are unknown and only a set P is known that contains both p_1 and p_2. In this case, the set of models is $\Theta = P \times P$, the model θ^* is any pair $p_1 = p_2$, and q_X^* is the strategy that for every signal x decides about the state y with equal probabilities $q_X^*(y = 1|x) = q_X^*(y = 2|x)$. This strategy is quite bad but no better strategy is possible. However, it is clear, at least intuitively, that the situation could be

improved if a learning sample was available. Minimax strategies are not able to take advantage of a learning sample, no matter how large. ■

Theorem 4 is a severe negative result. The result is quite transparent and easy to prove, especially, if one notices that the condition of the theorem states that the function $R_X : Q_X \times \Theta \to R$ has a saddle point and the pair (q_X^*, θ^*) is this point. Probably, the property was discovered many times before. In a special case, this property was noticed by H.Robbins when he wanted to derive the empirical Bayesian approach from the ideas of minimax decisions [16]. This clarity may have been the reason for the hasty and pessimistic conclusion that recognition with learning is something very different from recognition without learning, something beyond the statistical decision theory.

Theorems 3 and 4 show the state of today's knowledge of complex object recognition as well as a gap in this knowledge. The minimax approach is suited for recognizing hidden state based on the signal that depends both on the state and on the model. However, for some complex objects this approach ignores learning information that depends directly only on the model and does not depend on the state. Unlike minimax approach, maximum likelihood learning makes use of the learning sample. It allows to recognize complex objects almost as if the model of the object is known. However, it is possible only when the learning sample size can be increased almost infinitely. For limited samples and for some objects maximum likelihood learning is not Bayesian and is predominated by some other methods. Therefore, today's developers have to cope alone with the problem when the model of an object is not completely known but several few examples of object's behaviour (let us say, 4-5) are available.

Let us see now how the closest to optimal learning procedures behave. As shown by the following theorem similar to Theorem 4, these procedures are also useless for a certain class of complex objects.

Theorem 5. *Let for a complex object* $\langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle$ *a model* θ^* *and a strategy* q_X^* *exist such that*

$$q_X^* = \operatorname*{argmin}_{q_X \in Q_X} [R_X(q_X, \theta^*) - \min_{q'_X \in Q_X} R_X(q'_X, \theta^*)], \tag{28}$$

$$\theta^* = \operatorname*{argmax}_{\theta \in \Theta} [R_X(q_X^*, \theta) - \min_{q'_X \in Q_X} R_X(q'_X, \theta)]. \tag{29}$$

Then any source $\langle Z, \Theta, p : Z \times \Theta \to \mathbb{R} \rangle$ *of learning information and any learning procedure* g_Z: $Z \to Q$ *satisfy the inequality*

$$\max_{\theta \in \Theta} [R_Z(g_Z, \theta) - \min_{q'_X \in Q_X} R_X(q'_X, \theta)] \geq \max_{\theta \in \Theta} [R_X(q_X^*, \theta) - \min_{q'_X \in Q_X} R_X(q'_X, \theta)].$$

However, the consequences of this theorem for closest to optimal learning are not so destructive as those of Theorem 4 for minimax learning. In fact, conditions (28) and (29) imply that an overall optimal Bayesian strategy exists

for the recognized object, namely the strategy q_X^*. Really, it follows from (28) that

$$q_X^* = \operatorname*{argmin}_{q_X \in Q_X}[R_X(q_X, \theta^*) - \min_{q'_X \in Q_X} R_X(q'_X, \theta^*)] = \operatorname*{argmin}_{q_X \in Q_X} R_X(q_X, \theta^*),$$

$$R_X(q_X^*, \theta^*) - \min_{q'_X \in Q_X} R_X(q'_X, \theta^*) = 0.$$

Thus, a condition (29) obtains a form

$$0 \geq R_X(q_X^*, \theta) - \min_{q'_X \in Q_X} R_X(q'_X, \theta) \text{ for all } \theta \in \Theta.$$

Evidently,

$$R_X(q_X^*, \theta) - \min_{q'_X \in Q_X} R_X(q'_X, \theta) \geq 0 \text{ for all } \theta \in \Theta,$$

and so

$$R_X(q_X^*, \theta) = \min_{q'_X \in Q_X} R_X(q'_X, \theta) \text{ for all } \theta \in \Theta.$$

In this case, any learning approach is useless, not only closest to optimal one.

Just as maximum likelihood learning, closest to optimal learning is also asymptotically optimal in the sense of the following theorem.

Theorem 6. *Let $Z = (X \times Y)^n$ and the learning procedure g_n^*: $(X \times Y)^n \to Q$ be defined by*

$$g_n^* = \operatorname*{argmin}_{g_Z \in G_Z} \max_{\theta \in \Theta}[R_Z(g_Z, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta)].$$

Then

$$\lim_{n \to \infty} \max_{\theta \in \Theta}[R_Z(g_n^*, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta)] = 0. \tag{30}$$

Proof. We proved earlier (see (25)) that

$$\lim_{n \to \infty} [R_Z(g_n^{\mathrm{ML}}, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta)] = 0$$

for procedure (24) and every model θ. Consequently,

$$\lim_{n \to \infty} \max_{\theta \in \Theta}[R_Z(g_n^{\mathrm{ML}}, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta)] = 0 \tag{31}$$

because the number $|\Theta|$ of models is finite. Clearly,

$$\min_{g_Z \in G_Z} \max_{\theta \in \Theta}[R_Z(g_Z, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta)] \leq \max_{\theta \in \Theta}[R_Z(g_n^{\mathrm{ML}}, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta)]. \tag{32}$$

Equality (31) and inequality (32) imply (30). ■

Hence, the concept of closest to optimal learning fills the above-mentioned gap in complex object recognition. It embraces the whole range of learning sample sizes, from zero to infinity. Unlike minimax approach it ignores the learning sample only if an overall optimal strategy exists for the recognized object. Certainly, in this case no information that specifies the model is needed. Similarly to the risks of maximum likelihood learning, the risks of closest to optimal learning converge to the risks of the optimal strategy. Unlike maximum likelihood learning, closest to optimal learning is Bayesian and is predominated by no other learning procedure.

5 Developing closest to optimal learning procedures

We consider a complex object, presented with a quadruple $\langle X, Y, \Theta, p : X \times Y \times \Theta \to \mathbb{R} \rangle$, and a source of learning information, presented with a triple $\langle Z, \Theta, p : Z \times \Theta \to \mathbb{R} \rangle$. The following two theorems show the way of developing the closest to optimal learning procedure. Clearly, this procedure is predominated with no procedure and therefore it is Bayesian. The next theorem specifies a weight function, with respect to which it is Bayesian.

Theorem 7. *The closest to optimal learning procedure*

$$g_Z^* = \operatorname*{argmin}_{g_Z \in G_Z} \max_{\theta \in \Theta} \big[R_Z(g_Z, \theta) - \min_{q_X \in Q_X} R_X(q, \theta) \big]$$

is a Bayesian procedure with respect to weights $\tau^*(\theta), \theta \in \Theta$ that maximize the difference

$$\left[\min_{g_Z \in G_Z} \sum_{\theta \in \Theta} \tau(\theta) \cdot R_Z(g_Z, \theta) \right] - \left[\sum_{\theta \in \Theta} \tau(\theta) \cdot \min_{q_X \in Q_X} R_X(q_X, \theta) \right].$$

Proof. Let us define a function $F\colon T \times G_Z \to \mathbb{R}$ given by

$$F(\tau, g_Z) = \left[\sum_{\theta \in \Theta} \tau(\theta) \cdot R_Z(g_Z, \theta) \right] - \left[\sum_{\theta \in \Theta} \tau(\theta) \cdot \min_{q_X \in Q_X} R_X(q_X, \theta) \right]$$

that allows to express procedure g_Z^* and weight function τ^* in a form

$$g_Z^* = \operatorname*{argmin}_{g_Z \in G_Z} \max_{\tau \in T} F(\tau, g_Z), \quad \tau^* = \operatorname*{argmax}_{\tau \in T} \min_{g_Z \in G_Z} F(\tau, g_Z).$$

We have to prove that $g_Z^* = \operatorname*{argmin}_{g_Z \in G_Z} \sum_{\theta \in \Theta} \tau^*(\theta) R_Z(g_Z, \theta)$. The function $F\colon T \times G_Z \to \mathbb{R}$ is a linear function of g_Z for every fixed τ and a linear function of τ for every fixed g_Z. It is defined on the Cartesian product of two convex closed sets T and G_Z. Due to the well-known duality theorem [2, 3, 10] this function satisfies the equality

$$\min_{g_Z \in G_Z} \left[\max_{\tau \in T} F(\tau, g_Z) \right] = F(\tau^*, g_Z^*) = \max_{\tau \in T} \left[\min_{g_Z \in G_Z} F(\tau, g_Z) \right].$$

This implies the equality $F(\tau^*, g_Z^*) = \min\limits_{g_Z \in G_Z} F(\tau^*, g_Z)$ and the chain

$$g_Z^* = \operatorname*{argmin}_{g_Z \in G_Z} F(\tau^*, g_Z) = \operatorname*{argmin}_{g_Z \in G_Z} \left[\sum_{\theta \in \Theta} \tau^*(\theta) R_Z(g_Z, \theta) - \sum_{\theta \in \Theta} \tau^*(\theta) \min_{q_X \in Q_X} R(q, \theta) \right] =$$

$$= \operatorname*{argmin}_{g_Z \in G_Z} \sum_{\theta \in \Theta} \tau^*(\theta) R_Z(g_Z, \theta).$$

■

For each $g_Z \in G_Z$, the function $F(\tau, g_Z)$ depends linearly on weights $\tau(\theta)$, $\theta \in \Theta$. Therefore, evidently, the function $\Phi(\tau) = \min_{g_Z \in G_Z} F(\tau, g_Z)$ is a concave function of these weights. Nevertheless, we deduce this fact directly from the definition of concavity and obtain an explicit expression for the supergradient of the function Φ, which is necessary for its further maximization.

Theorem 8. *The function* $\Phi\colon \mathbb{R}^\Theta \to \mathbb{R}$ *with values*

$$\Phi(\tau) = \left[\min_{g_Z \in G_Z} \sum_{\theta \in \Theta} \tau(\theta) \cdot R_Z(g_Z, \theta) \right] - \left[\sum_{\theta \in \Theta} \tau(\theta) \cdot \min_{q_X \in Q_X} R_X(q_X, \theta) \right]$$

is a concave function of weights $\tau(\theta)$, $\theta \in \Theta$.

Proof. We show that for every point $\tau^* \in \mathbb{R}^\Theta$ some linear function $L\colon \mathbb{R}^\Theta \to \mathbb{R}$ exists that satisfies the inequality

$$L(\tau - \tau^*) \geq \Phi(\tau) - \Phi(\tau^*) \tag{33}$$

for every $\tau \in \mathbb{R}^\Theta$. Note that the gradient of this linear function is a supergradient of the function Φ at the point τ^*. One of possible linear functions that satisfy (33) is the function with values

$$L(\tau) = \sum_{\theta \in \Theta} \tau(\theta) \left[R_Z(g_Z^*, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta) \right] \quad \text{where} \quad g_Z^* = \operatorname*{argmin}_{g_Z \in G_Z} \sum_{\theta \in \Theta} \tau^*(\theta) R_Z(g_Z, \theta).$$

Indeed,

$$L(\tau - \tau^*) =$$

$$= \sum_{\theta \in \Theta} \tau(\theta) \left[R_Z(g_Z^*, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta) \right] - \sum_{\theta \in \Theta} \tau^*(\theta) \left[R_Z(g_Z^*, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta) \right] \geq$$

$$\geq \min_{g_Z \in G_Z} \sum_{\theta \in \Theta} \tau(\theta) \left[R_Z(g_Z, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta) \right] -$$

$$- \min_{g_Z \in G_Z} \sum_{\theta \in \Theta} \tau^*(\theta) \left[R_Z(g_Z, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta) \right] = \Phi(\tau) - \Phi(\tau^*).$$

■

Due to concavity of the function Φ, it can be maximized using well-known methods of convex optimization, such as supergradient hill-climbing [21]. To do that, an auxiliary procedure to project a point $\tau \in \mathbb{R}^{\Theta}$ onto the set T is needed. Strictly speaking, this procedure finds the point $\operatorname{argmin}_{\tau' \in T}(\tau' - \tau)^2$ for a given point $\tau \in \mathbb{R}^{\Theta}$. In our case the procedure is as follows:

1. for all $\theta \in \Theta$ do $\tau'(\theta) = \tau(\theta)$;
2. for all $\theta \in \Theta$ do $\tau'(\theta) = \tau'(\theta) - (\sum_{\theta \in \Theta} \tau'(\theta) - 1) \cdot |\Theta|^{-1}$;
3. if $\tau'(\theta) \geq 0$ for all $\theta \in \Theta$ then $STOP$;
4. for all $\theta \in \Theta$ do $\tau'(\theta) = \max\{\tau'(\theta), 0\}$;
5. go to 2;

Next we present a technique to compute the closest to optimal learning procedure. We use the word 'technique' rather than 'algorithm', to emphasize that some of its steps are far from trivial and have to be developed specifically for every task at hand. The technique can be seen as an algorithm only in very simple cases, such as the examples given in Section 6.

Initially, the following input data have to be given:

1. sets X, Y, Θ and probabilities $p(x, y; \theta)$, $x \in X$, $y \in Y$, $\theta \in \Theta$, which define the complex object to be recognized;
2. set Z and probabilities $p(z; \theta)$, $z \in Z$, $\theta \in \Theta$, which define the source of learning information;
3. values of the loss function $w(y, y')$ for every $y \in Y, y' \in Y$
4. numbers $\min_{q_X \in Q_X} R_X(q_X, \theta)$ for every $\theta \in \Theta$;
5. sequence of numbers γ_i, $i = 1, 2, \ldots$, such that $\lim_{i \to \infty} \gamma_i = 0$ and $\lim_{n \to \infty} \sum_{i=1}^{n} \gamma_i = \infty$;
6. required accuracy $\varepsilon > 0$;
7. initial weights $\tau(\theta)$, $\theta \in \Theta$; for example, $\tau(\theta) = |\Theta|^{-1}$;
8. initial values of numbers S and s, which will keep an upper and lower bound of the quality of the learning procedure; for example $S = \infty$ and $s = 0$.

The i-th iteration of the technique proceeds as follows:

1. construct a Bayesian procedure g_Z for the current weight function τ;
2. for every model $\theta \in \Theta$ calculate numbers $\Delta(\theta) = R_Z(g_Z, \theta) - \min_{q_X \in Q_X} R_X(q_X, \theta)$; they form the gradient of function Φ at the current point τ;
3. calculate $\Delta = \max_{\theta} \Delta(\theta)$; it is the quality of the current procedure g_Z;
4. calculate $\bar{\Delta} = \sum_{\theta} \tau(\theta) \Delta(\theta)$; it is the value of $\Phi(\tau)$ at the current point τ;
5. update the bounds as $S = \min\{S, \Delta\}$ and $s = \max\{s, \bar{\Delta}\}$;
6. if $S - s < \varepsilon$ then $STOP$;

7. set the weights $\tau(\theta)$, $\theta \in \Theta$, to the projection of the weights $\tau(\theta) + \gamma_i \Delta(\theta)$ on T;

Supergradient optimization, which iteratively updates the initial weights, is performed only in steps 2, 3 and 7. The other steps calculate the data used in the stopping condition, which, as will be shown, guarantees the required accuracy.

Let S and s be the bounds after termination of the algorithm. It means that a learning procedure g'_Z with the quality

$$\max_{\theta\in\Theta}[R_Z(g'_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)] = S \tag{34}$$

has been obtained at some iteration. It also means that a weight function τ' with

$$\Phi(\tau') = \left[\min_{g_Z\in G_Z}\sum_{\theta\in\Theta}\tau'(\theta)\cdot R_Z(g_Z,\theta)\right] - \left[\sum_{\theta\in\Theta}\tau'(\theta)\cdot \min_{q_X\in Q_X} R_X(q_X,\theta)\right] = s \tag{35}$$

has been obtained at some (possibly different) iteration. As shown by the following theorem, the procedure g'_Z and the weight function τ' have the property that the difference between the quality of the obtained procedure and the best procedure is not greater than ε.

Theorem 9. *The learning procedure g'_Z satisfies the inequality*

$$\max_{\theta\in\Theta}[R_Z(g'_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)] - \min_{g_Z\in G_Z}\max_{\theta\in\Theta}[R_Z(g_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)] < \varepsilon.$$

Proof. As before, we consider the function

$$F(\tau, g_Z) = \sum_{\theta\in\Theta}\tau(\theta)\big[R_Z(g_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)\big].$$

By (34) and (35), for the procedure g'_X and the weight function τ' we have the chain

$$\begin{aligned}
s &= \min_{g_Z\in G_Z}\sum_{\theta\in\Theta}\tau'(\theta)\big[R_Z(g_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)\big] = \min_{g_Z\in G_Z} F(\tau', g_Z)\\
&\le \max_{\tau\in T}\min_{g_Z\in G_Z} F(\tau, g_Z) = \min_{g_Z\in G_Z}\max_{\tau\in T} F(\tau, g_Z)\\
&= \min_{g_Z\in G_Z}\max_{\tau\in T}\sum_{\theta\in\Theta}\tau(\theta)\big[R_Z(g_Z,\theta) - \min_{q_X\in Q} R_X(q_X,\theta)\big]\\
&= \min_{g_Z\in G_Z}\max_{\theta\in\Theta}[R_Z(g_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)]\\
&\le \max_{\theta\in\Theta}[R_Z(g'_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)\big] = S.
\end{aligned}$$

From this chain and the stopping condition $S - s < \varepsilon$, it follows that

$$\max_{\theta\in\Theta}[R_Z(g'_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)] - \min_{g_Z\in G_Z}\max_{\theta\in\Theta}[R_Z(g_Z,\theta) - \min_{q_X\in Q_X} R_X(q_X,\theta)] < \varepsilon.$$

■

6 Examples

For several simplest cases, the closest to optimal learning procedure g^0 is compared with the maximum likelihood learning procedure g^{ML} and with another learning procedures g^H (to be defined later). In all the examples, we have an object with sets $X = \mathbb{R}$, $Y = \{1, 2\}$ and probability density distribution of the form

$$p(x, y; \theta) = p_y \cdot (\sqrt{2\pi})^{-1} \cdot \mathrm{e}^{-\frac{1}{2}(x-\mu_y)^2}$$

with some unknown parameters specified later. Such object is a complex object in a sense of Definition 2 and the compared learning procedures for it can be implemented almost exactly. We use the loss function (5). The required accuracy for constructing closest to optimal procedures is $\varepsilon = 0.01$.

Example 8. [18] Let $\mu_1 = 1$, $\mu_2 = (-1)$, and only the a priori probabilities p_y, $y \in \{1, 2\}$, be unknown. Thus, the set of models is $\Theta = \{\theta \mid 0 \le \theta \le 1\}$ and the a priori probabilities p_y of the states are $p_1 = \theta$, $p_2 = 1 - \theta$. This is the example used by H.Robbins to explain his idea of empirical Bayesian approach. We simplify the example even further by assuming that the learning information is not a signal sample $(x_1, x_2, \ldots, x_n)$ but a state sample $(y_1, y_2, \ldots, y_n)$. To use the technique from Section 5 the set Θ is discretized with a step 0.05.

Figure 2 (page 12) shows how the probabilities $\min_{q \in Q} R(q, \theta)$ and $R(q^{\text{minmax}}, \theta)$ depend on model θ. Note that they do not depend on the sample size n because they do not depend on the sample at all. Figure 4 shows how the risks $R_G(g^{ML}, \theta)$ and $R_G(g^0, \theta)$ depend on the model θ for the sample sizes $n = 1, 2, 4$.

It is striking how high the risk $R_G(g^{ML}, \theta)$ is for some models. Of course, one cannot expect the risk to be too low because the learning samples are very small and thus they bring little information about the true model. Nevertheless, however small the amount of information, it is non-zero and it should be used to improve subsequent recognition, at least to some extent. But the example shows that it is much better to ignore this information than to use it with maximum likelihood learning. The example illustrates quite transparently the main drawback of maximum likelihood learning. ■

Example 9. For the same complex object as in Example 8, let the learning information be a sequence $(x_1, x_2, \ldots, x_n)$ rather than $(y_1, y_2, \ldots, y_n)$. This is exactly the case considered by H.Robbins. The risk $R_G(g^0, \theta)$ of closest to optimal learning is compared with the risk $R_G(g^{ML}, \theta)$ of maximum likelihood learning. Even for this simple complex object, it is not easy to find the maximum likelihood model

$$\theta_1 = \operatorname*{argmax}_{\theta \in \Theta} \sum_{i=1}^{n} \log\Big[\theta \cdot \mathrm{e}^{-\frac{1}{2}(x_i-1)^2} + (1-\theta) \cdot \mathrm{e}^{-\frac{1}{2}(x_i+1)^2}\Big].$$

Therefore, we also consider the heuristic procedure (10) by H.Robbins, which is denoted by g^H. The procedure is not based on the maximum likelihood

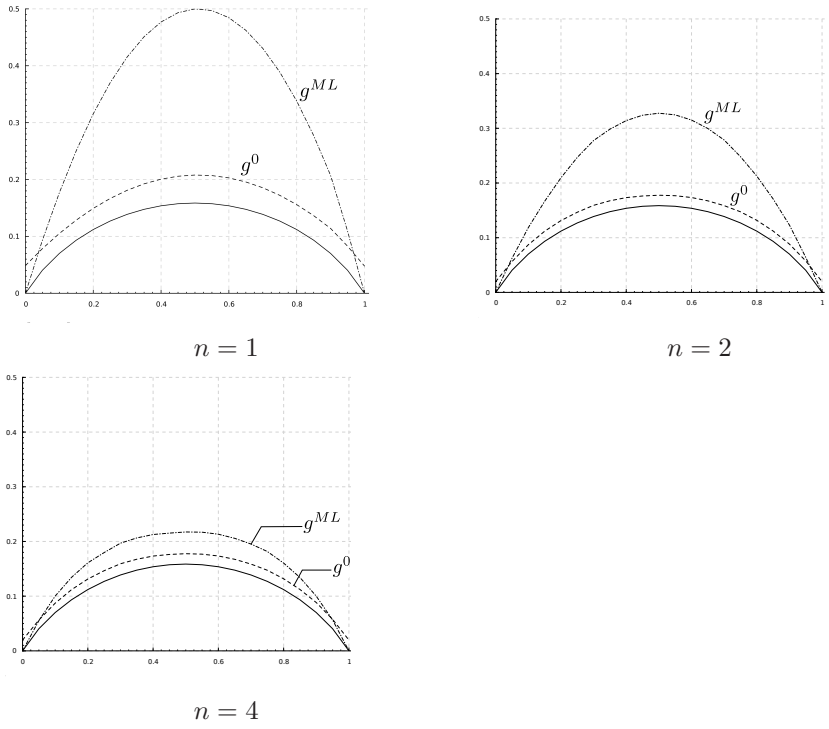

Figure 4: EXAMPLE 8. The dependency of $R_G(g^0, \theta)$ and $R_G(g^{ML}, \theta)$ on θ for sample sizes $n = 1, 2, 4$. The solid curve shows risk $\min_{q \in Q} R(q, \theta)$.

estimate θ_1 but on the consistent estimate $\theta_2 = \frac{1}{2n} \sum_{i=1}^{n} x_i + \frac{1}{2}$, which can be easily calculated. Figure 5 shows how the risks $R_G(g^0, \theta)$, $R_G(g^{ML}, \theta)$ and $R_G(g^H, \theta)$ depend on the model θ for the sample sizes $n = 1, 2, 4$. ■

Example 10. Let $X = \mathbb{R}$, $Y = \{1, 2\}$, $p_1 = p_2 = 0.5$, and $\mu_1 = 0$. The expected value $\mu_2 = \theta$ is unknown, it is only known that it belongs to the set $\Theta = \{-6, -5.8, -5.6, \ldots, 5.8, 6\}$ of 61 possible values. Thus,

$$p(x, y = 1; \theta) = 0.5 \cdot (\sqrt{2\pi})^{-1} \cdot \mathrm{e}^{-\frac{1}{2}x^2},$$
$$p(x, y = 2; \theta) = 0.5 \cdot (\sqrt{2\pi})^{-1} \cdot \mathrm{e}^{-\frac{1}{2}(x-\theta)^2}.$$

The learning information is a sequence $(x_1, x_2, \ldots, x_n)$ of random signals generated by the object in state $y = 2$. We did not observe any significant difference between maximum likelihood learning and closest to optimal learning. ■

Example 11. Here we have the same object as in the previous example but

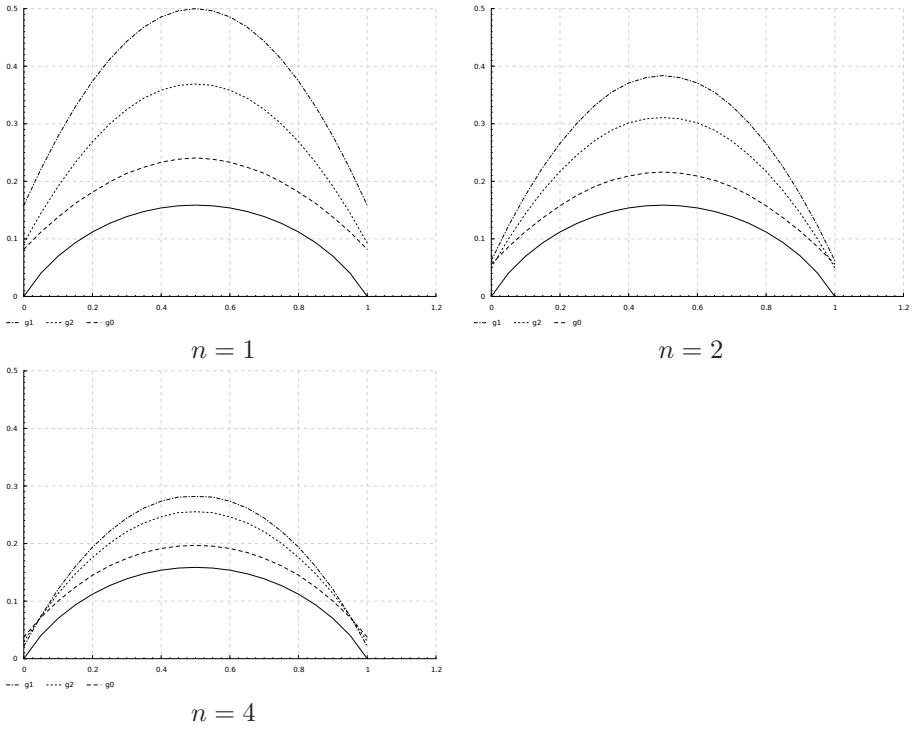

Figure 5: EXAMPLE 9. The dependency of $R_G(g^0, \theta)$, $R_G(g^{ML}, \theta)$, $R_G(g^H, \theta)$ on θ for sample sizes $n = 1, 2, 4$. The solid curve shows the risk $\min_{q \in Q} R(q, \theta)$.

the learning information is more complicated. It is a sequence $(x_1, x_2, \ldots, x_n)$ of signals generated by the object with unknown states. In this case, the maximum likelihood estimate

$$\theta_1 = \operatorname*{argmax}_{\theta \in \Theta} \sum_{i=1}^{n} \log \left[\mathrm{e}^{-\frac{1}{2} x_i^2} + \mathrm{e}^{-\frac{1}{2}(x_i - \theta)^2} \right]$$

is not easy to find. Therefore, besides the maximum likelihood estimate we also consider the consistent estimate $\theta_2 = \frac{2}{n} \sum\limits_{i=1}^{n} x_i$ along with the corresponding learning procedure g^H. Figure 6 shows how the risks $R_G(g^0, \theta)$, $R_G(g^{ML}, \theta)$, $R_G(g^H, \theta)$ depend on the model θ for the sample sizes $n = 1, 2, 4$. One can see that neither the maximal likelihood procedure g^{ML} nor the heuristic procedure g^H dominates the other. However, the closest to optimal procedure g^0 dominates both g^{ML} and g^H, especially for small sample sizes. We have not observed any significant difference between the compared procedures for the learning sample sizes larger than 10.

The case of an empty learning sample ($n = 0$) is of a particular interest. In this case the maximum-likelihood learning problem cannot even be formulated

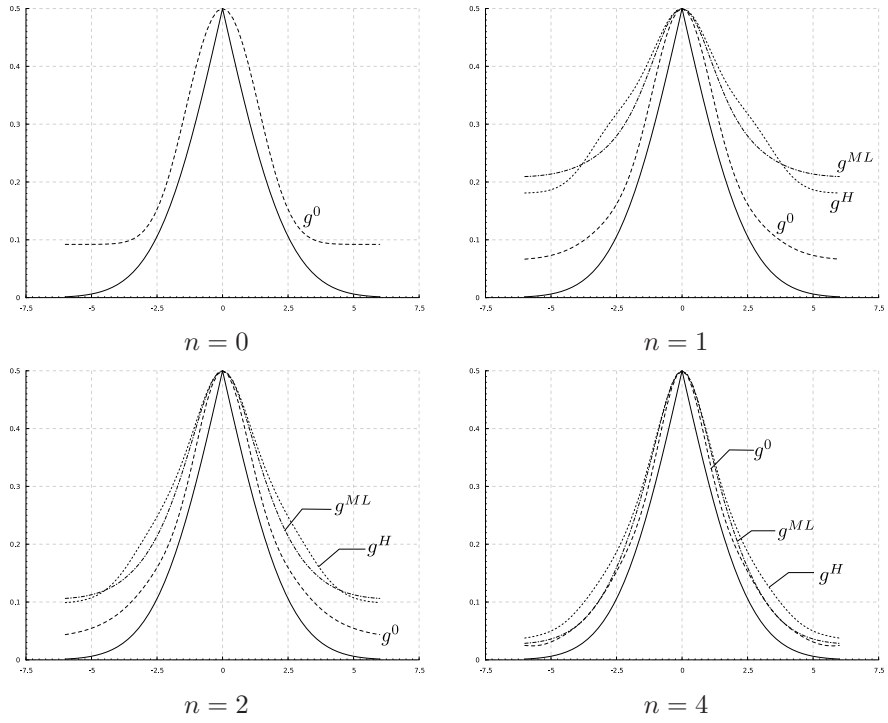

Figure 6: EXAMPLE 11. Dependency of the risks $R_G(g^0, \theta)$, $R_G(g^{ML}, \theta)$, $R_G(g^H, \theta)$ on the model θ for sample sizes $n = 0, 1, 2, 4$. The solid curve shows the risk $\min_{q \in Q} R(q, \theta)$.

because any model estimate is meaningless when empirical data are absent. The minimax requirement for recognition strategy may be formally stated but it result in a strategy that makes a wrong decision with probability 0.5 for each true model. As for the closest to optimal requirement, it can be formulated for this case as well and results in a quite reasonable strategy. Figure 6 shows how the risk $R_G(g^0, \theta)$ depends on θ when the learning sample is absent at all.

■

7 Results and corollaries

A common feature of known methods of recognition with learning is that a single model is selected from a given model set. This selected model 'best' matches the learning sample, usually it is a consistent estimate of the model. Nevertheless, only a single model is selected from the set of all possible models. Then the

recognition strategy is constructed as if the selected model was the true one. Due to consistency of the model estimate, such methods are acceptable when the learning sample is large enough. However, the methods give no guarantees for subsequent recognition when the learning sample has limited size, especially when the size is obviously insufficient to determine the model unambiguously. The drawback is well-known as the small sample problem and forms an obvious gap in today's knowledge of recognition with uncertain statistical model. Strictly speaking, clarity is achieved only in two extreme cases. If no learning sample is available, then it is appropriate to use the minimax approach. If the sample is large enough, then it is possible to use known methods of recognition with learning. If the situation falls somewhere inbetween, the application developer faces the question whether he should consider the sample large enough and use known methods of learning, or consider the sample too short and ignore it.

The concept of closest to optimal strategy allows to consider the recognition with learning from a point of view that differs from the conventional one. It answers some questions that were unanswered up to now. However, new questions arise that are not visible from the conventional viewpoint. Moreover, it turns out that the third and many other points of view on the problem are possible.

1. Both recognition without learning and recognition with learning can be formalized in a unified framework as developing a closest to optimal recognition strategy. The formalization covers the whole range of learning sample sizes, including the zero size.

2. The concept of closest to optimal strategy enables to formalize the main idea of empirical Bayesian approach of H.Robbins, namely that simultaneous recognition of multiple objects can be performed much better than recognition of each object independently. The problem of compound object recognition can be formulated as finding a closest to optimal strategy.

3. The uniform formulation of the recognition problem with respect to the sample size allows to achieve a satisfactory clarity in the small sample problem. Each learning sample, however small, is more useful than its absence. It is not the size of the learning sample that determines its sufficiency for subsequent recognition, it is not even the learning sample itself. Any sample, however small, can be sufficient for recognizing some objects, as well as even a very large sample can be insufficient for recognizing some other objects. It is the pair "sample–signal" that is sufficient or insufficient for recognition. It is not necessary to amend the developed approach in order to distinguish between sufficient and insufficient pairs "sample–signal". The discrimination can be achieved with tools that are already present in the approach. The decision about current object state has to take values from the set $Y \cup \{\sharp\}$ where $\sharp$ is the decision that the given pair "learning sample–signal" is insufficient to recognize the hidden state with a small enough risk. Thus, if the small sample problem is correctly formulated, it disappears as an independent problem. It is subsumed by the correct formulation of the recognition problem and inside this more general problem it is solved more subtly, more exactly.

4. Similarly, recognition with learning becomes a special case of recognition without learning. In this particular case, the data available for recognition

consist of two parts. The first part depends both on the hidden object state and on the unknown model. The second part directly depends on the model and does not depend on the state when the model is fixed.

5. To find the closest to optimal strategy for a realistic application, it is necessary to overcome fundamental computational difficulties. This is nothing new in the pattern recognition practice. Maximum likelihood estimation or empirical risk minimization are also far from trivial even in very simple realistic situations. In practice, they need to be implemented with certain acceptable simplifications. Similar difficulties can be expected for the developed approach.

6. There is only one universal statement in the presented approach. It is the dichotomy between Bayesian and improper strategies. In contrast, closest to optimal strategy can be defined in many other reasonable ways. According to our definition, a strategy q_X^* is closest to optimal if it satisfies the conditions

$$R_X(q_X^*, \theta) - \min_{q'_X \in Q_X} R_X(q'_X, \theta) \le c, \quad \theta \in \Theta,$$

with the minimal value of c. If the risks $R_X(q_X, \theta)$ are not negative, the conditions

$$R_X(q_X^*, \theta) \le c \cdot \min_{q'_X \in Q_X} R_X(q'_X, \theta), \quad \theta \in \Theta,$$

are no less reasonable. Moreover, it is not obligatory to use the optimal risks $\min\limits_{q \in Q_X} R_X(q, \theta)$ in these conditions. Any other values $R^*(\theta)$ can be used that express the desired parameters of the recognition system under development. However, independently of the possibly modified requirements, the developed strategy has to be Bayesian and its search is reduced to looking for an appropriated weight function over the set of possible models.

The article answers the questions that were formulated in [19, page 272] several years ago. The questions gave an opportunity to look at the problem from an unconventional point of view. Now, new questions arise and one can see that in spite of the long history of recognition with learning, a much longer way is ahead.

We are grateful to T.Werner, J.Matas and B.Flach for fruitful discussions, critical remarks and valuable advice.

References

[1] Rocío Alaiz-Rodríguez, Alicia Guerrero-Curieses, and Jesús Cid-Sueiro. Minimax regret classifier for imprecise class distributions. *J. Mach. Learn. Res.*, 8:103–130, May 2007.

[2] J.M. Borwein and A.S. Lewis. *Convex Analysis and Nonlinear Optimization.* Springer Verlag, 2000.

[3] S. Boyd and L. Vandenberghe. *Convex Optimization.* Cambrige University Press, 2004.

[4] George Casella. An introduction to empirical bayes data analysis. *The American Statistician*, 39(2):83–87, 2009.

[5] H. Chernoff and H.C.L.E. Moses. *Elementary Decision Theory*. Wiley series in probability and mathematical statistics: Applied probability and statistics. Dover Publ., 1959.

[6] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977.

[7] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. Wiley, 2000.

[8] M. Ehrgott. *Multicriteria Optimization*. Lecture notes in economics and mathematical systems. Springer, 2005.

[9] Arthur M. Geoffrion. Proper Efficiency and the Theory of Vector Maximization. *Journal of Mathematics, Analysis and Applications*, 22(3):618–630, June 1968.

[10] J.-B. Hiriart-Urruty and C. Lemarechal. *Fundamentals of Convex Analysis*. Springer Verlag, 2002.

[11] M.G. Kendall and A. Stuart. *The Advanced Theory of Statistics: In Three Volumes*. Number 2 in The Advanced Theory of Statistics. Griffin, 1979.

[12] Josef Kittler, Mohamad Hatef, Robert P. W. Duin, and Jiri Matas. On combining classifiers. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 20:226–239, 1998.

[13] E.L.A. Lehmann. *Testing Statistical Hypotheses*. Springer Texts in Statistics. Springer Verlag, 1986.

[14] H.DeGroot Morris. *Optimal Statistical Decisions*. McGraw-Hill Company, 1970.

[15] J. Neyman. Two Breakthroughs in the Theory of Statistical Decision Making. *Revue de l'Institut International de Statistique / Review of the International Statistical Institute*, 30(1):11–27, 1962.

[16] Herbert Robbins. Asymptotically Subminimax Solutions of Compound Statistical Decision Problems. In Jerzy Neyman, editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 131–148. University of California Press, 1951.

[17] Herbert Robbins. An empirical bayes approach to statistics. *Proc. Third Berkeley Symp. on Math. Statist. and Prob.*, 1:157–163, 1956.

[18] M.I. Schlesinger and A. V. Bondarenko. On pattern recognition learning problem formulation. (in russian). *Control Systems and Computers*, 2:4–19, 2009.

[19] M.I. Schlesinger and V. Hlavac. *Ten Lectures on Statistical and Structural Pattern Recognition*. Computational Imaging and Vision Series. Kluwer Academic Pub, 2002.

[20] M.I. Shlezinger. The interaction of learning and self-organization in pattern recognition. *Kibernetika*, 2:81–88, 1968.

[21] N.Z. Shor. *Nondifferentiable Optimization and Polynomial Problems*. Nonconvex Optimization and Its Applications. Springer, 1998.

[22] Andrew R. Webb. *Statistical Pattern Recognition*. Wiley, 2002.

[23] M. Zeleny. *Multiple criteria decision making*. McGraw-Hill series in quantitative methods for management. McGraw-Hill, 1982.